A multi-paradigm language for reactive synthesis

Ioannis Filippidis Richard M. Murray
{ifilippi, murray}@caltech.edu
Control and Dynamical Systems,
California Institute of Technology,
Pasadena CA 91125, USA

Gerard J. Holzmann
gerard.j.holzmann@jpl.nasa.gov
Laboratory for Reliable Software,
Jet Propulsion Lab, Caltech,
Pasadena CA 91109, USA

This paper proposes a language for describing reactive synthesis problems that integrates imperative and declarative elements. The semantics is defined in terms of two-player turn-based infinite games with full information. Currently, synthesis tools accept linear temporal logic (LTL) as input, but this description is less structured and does not facilitate the expression of sequential constraints. This motivates the use of a structured programming language to specify synthesis problems. Transition systems and guarded commands serve as imperative constructs, expressed in a syntax based on that of the modeling language PROMELA. The syntax allows defining which player controls data and control flow, and separating a program into assumptions and guarantees. These notions are necessary for input to game solvers. The integration of imperative and declarative paradigms allows using the paradigm that is most appropriate for expressing each requirement. The declarative part is expressed in the LTL fragment of generalized reactivity(1), which admits efficient synthesis algorithms, extended with past LTL. The implementation translates PROMELA to input for the SLUGS synthesizer and is written in PYTHON. The AMBA AHB bus case study is revisited and synthesized efficiently, identifying the need to reorder binary decision diagrams during strategy construction, in order to prevent the exponential blowup observed in previous work.

1 Introduction

Over the past three decades, system formal verification has aided design and become practical for industrial application. In the past decade, synthesis of systems from specifications has seen significant development [60, 100], partially owing to the discovery of temporal logic fragments that admit efficient synthesis algorithms [83, 20, 30, 8]. Applications range from protocol synthesis for hardware circuits [20], to correct-by-construction controllers for hybrid systems [57, 56, 101].

Many languages and tools have been developed for modeling and model checking systems. Unlike verification using model checking, the tools for synthesis have been developed much more recently. One reason is that centralized synthesis from linear temporal logic (LTL) [85] has doubly exponential complexity in the length of the specification formula [89], a result that did not encourage further development initially.

Currently, LTL is the language used for describing specifications as input to synthesis tools. There are many benefits in using a logic for synthesis tasks, its declarative nature being a major one, because it allows expressing individual requirements separately, and in a precise way. It also makes explicit the implicit conventions present in programming languages [61]. Another aspect of synthesis problems that makes declarative descriptions appropriate is that we want to describe as large a set of possible designs as possible, in order to avoid overconstraining the search space.

However, not all specifications are best described declaratively. There exist synthesis problems whose description involves graph-like structures that are cumbersome for humans to write in logic. Robotics problems typically involve graph constraints that originate from possible physical configurations. For example, considering a wheeled robot, its physical motion is modeled by possible transitions

Černý, Kuncak, Madhusudan (Eds.): Fourth Workshop on
Synthesis (SYNT 2015)
EPTCS 202, 2016, pp. 73–97, doi:10.4204/EPTCS.202.6

© I. Filippidis, R.M. Murray, G.J. Holzmann

that avoid collisions with other objects, whereas an objective to patrol between two locations can more appropriately be described with a temporal logic formula. Properties that specify sequential behavior also lead to graph-like structures, and require use of auxiliary variables that serve as memory. Expressing sequential composition in logic leads to long, unstructured formulas that deemphasize the specifier's intent. The resulting specifications are difficult to maintain, and writing them is error-prone. In addition, the specifier may need to explicitly write clauses that constrain variables to remain unchanged, in order to maintain imperative state. This leads to longer formulas in which the intent behind individual clauses is less readable.

Another motivation relates to the temporal logic hierarchy [71, 92]. Synthesis from LTL has time complexity polynomial in the state space size, and doubly exponential in the size of the formula. In contrast, algorithms with linear time complexity in the size of the formula are known for the fragment of generalized reactivity of rank one, known as GR(1).

In the automata hierarchy, the GR(1) fragment corresponds to an implication of deterministic Büchi automata (BAs) [83, 92, 78, 93]. The consequent requires some system behavior, provided that the environment satisfies the antecedent of the implication, as described in Section 2.3. Deterministic automata can describe recurrence properties ($\square\diamond$), but not persistence ($\diamond\square$). Intuitively, the behavior of variables uniquely determines the associated behavior of a deterministic BA. This drops the complexity of synthesis, because the algorithm does not have to keep track of branching in the automata that is not recorded in the problem's variable.

A large subset of properties that are of practical interest in industrial applications [28, 69, 20] can be expressed in GR(1). There do exist properties that cannot be represented by deterministic Büchi automata, e.g., persistence $\diamond\square p$. Of these properties, those with Rabin rank equal to one are still amenable to polynomial time algorithms (by solving parity games) [30]. Higher Rabin ranks are not expected to admit polynomial time solution, unless P = NP [30]. This motivates formulating the required properties in GR(1), which corresponds to Streett properties with rank one. The winning set for a Streett objective of rank one can be computed with the same time complexity as that for a Rabin objective of rank one. Therefore, properties in the lower Rabin ranks are known to be at least as hard to synthesize as GR(1). This motivates formulating the required properties in GR(1), which trades off expressive power for computational efficiency.

Translating properties to deterministic automata can be done automatically, but may lead to more expensive synthesis problems than manually written properties, as reported in [78]. So the ability to write deterministic automata directly in a structured and readable language avoids the need for automated translation, and allows fine tuning them, based on the specifier's understanding of the problem. The trade-off is that the translation has to be performed by a human.

Another reason why specifications are not always purely declarative is that in many cases we want to synthesize a system using *existing* components. In other words, we already have a *partial* model, which describes the possible behavior of components that already exist, e.g., because we purchased them off the self, to interface them with the part of the system that we are synthesizing. We *declare* to our synthesis tool what properties the controller under design should satisfy with respect to this model. This restricts what the system should achieve using these components, but not how exactly that will be achieved. So the partial model is best described imperatively, whereas the goal declaratively, using temporal logic.

Educationally, the transition for students from a general purpose programming language like PYTHON or `C`, directly to temporal logic constitutes a significant leap. Using a multiparadigm language can make this transition smoother.

This work proposes a language that can describe synthesis problems for open systems that react to an adversarial environment. The syntax is derived from that of PROMELA, whereas the semantics interprets

it as a two-player turn-based game of infinite duration. Both synchronous and asynchronously scheduled centralized systems with full information can be synthesized. In Section 2, we review temporal logic and relevant notions about two-player games. The presence of two players requires declaring who controls each variable (Section 3.1), as well as the data flow, and control flow in transition systems (Section 3.2). In addition, the specification needs to be partitioned into assumptions about the environment, and guarantees that the system must satisfy (Section 2, Section 3.2). The integration of declarative and imperative semantics is obtained by defining imperative variables (Section 3.1), deconstraining, and executability of actions (Section 3.3). In order to be synthesized, the program is translated to temporal logic, as described in Section 4. In Section 5, we discuss the implementation, and in Section 6 significant improvements in the AMBA case study [20] that were possible by merging fairness requirements into a single Büchi automaton. Relevant work is collected in Section 7, and conclusions in Section 8.

2 Preliminaries

2.1 Linear temporal logic

Linear temporal logic with past is an extension of Boolean logic used to reason about temporal modalities over sequences. The temporal operators next $\bigcirc$, previous $\ominus$, until $\mathscr{U}$, and since $\mathscr{S}$ suffice to define the other operators [85, 10]. Let AP be a set of variable symbols p that can take values over $\mathbb{B} \triangleq \{\bot, \top\}$. A model of an LTL formula is a sequence of variable valuations called a *word* $w : \mathbb{N} \to \mathbb{B}^{AP}$. A well-formed formula is inductively defined by $\varphi ::= p|\neg\varphi|p \wedge p|\bigcirc\varphi|\varphi\,\mathscr{U}\,\varphi|\ominus\varphi|\varphi\,\mathscr{S}\,\varphi$. A formula φ is evaluated over a word w at a time $i \geq 0$, and $w,i \models \varphi$ denotes that φ holds at position i of word w. Formula $\bigcirc\varphi$ holds at position i if φ holds at position $i+1$, $\varphi\,\mathscr{U}\,\psi$ holds at i if there exists a time $j \geq i$ such that $w,j \models \psi$ and for all $i \leq k < j$, it is $w,k \models \varphi$. The operator $\diamondsuit p \triangleq \top\,\mathscr{U}\,p$ requires that p be eventually true, and the operator $\square p \triangleq \neg\diamondsuit\neg p$ requires that p be true over the whole word. The past fragment of LTL extends it with the previous and since operators, $\ominus, \mathscr{S}$ respectively [66, 70, 54]. Formula $\ominus\varphi$ holds at i iff $i > 0$ and $w,i-1 \models \varphi$, and formula $\varphi\,\mathscr{S}\,\psi$ holds at i iff there exists a time j with $0 \leq j \leq i$ such that $w,j \models \psi$, and for all k such that $j < k \leq i$ it is $w,k \models \varphi$. The *weak previous* operator $\circledcirc$ is defined as $\circledcirc\varphi \triangleq \neg\ominus\neg\varphi$, once $\diamondsuit$ as $\diamondsuit\varphi \triangleq \top\,\mathscr{S}\,\varphi$, and historically $\boxminus$ as $\boxminus\varphi \triangleq \neg\diamondsuit\neg\varphi$. Past LTL is implemented using temporal testers [54].

2.2 Turn-based games

In many applications, we are interested in designing a system that does not have full control over the behavior of all variables that are used to model the situation. Some problem variables represent the behavior of other entities, usually collectively referred to as the environment. The system reads these *input* variables and reacts by writing to *output* variables that it controls, continuing indefinitely. Such a system is called *open* [6, 84], to distinguish it from closed systems that have no inputs, and so full control.

The synthesis of an open system can be formulated as a two-player adversarial game of infinite duration [96]. The two players in the game are usually called the protagonist (system) and antagonist (environment). We control the protagonist, but not the antagonist. If the players move in turns, then the game is called *alternating*. Each pair of consecutive moves by the two players is called a *turn* of the game. In each turn, player 0 moves first, without knowing how player 1 will choose to move in that turn of the game. Then player 1 moves, knowing how player 0 moved in that turn. Depending on which player we control, there are two types of game. If the protagonist is player 1, then the game is called

Mealy, otherwise *Moore* [75, 76]. Due to the difference in knowledge about the opponent's next move between the two flavors of game, more specifications are realizable in a Mealy game, than in a Moore game. There exist solvers for both Moore and Mealy games. Here we will consider Mealy games only.

2.3 Games in logic

Temporal logic can be used to describe both the possible moves in a game (the *arena* or *game graph*), as well as the winning condition. Let $\mathscr{X}$ and $\mathscr{Y}$ be two sets of propositional variables, controlled by the environment and system, respectively. Let $\mathscr{X}'$ and $\mathscr{Y}'$ denote primed variables, where x' represents the next value $\bigcirc x$ of variable x. We abuse notation by using primed variables inside temporal formulae.

Synthesis from LTL specifications is in 2EXPTIME [87, 89], motivating the search for fragments that admit more efficient synthesis algorithms. Generalized reactivity of index one, abbreviated as GR(1), is a fragment of LTL that admits synthesis algorithms of time complexity polynomial in the size of the state space [20]. GR(1) [50, 88, 16, 67, 32] is used in the following, but the results can be adapted to larger fragments of LTL, provided that another synthesizer be used [49, 31, 29, 21, 33].

The possible moves in a Mealy game can be specified by initial and transition conditions that constrain the environment and system. Initial conditions are described by propositional formulae over $\mathscr{X} \cup \mathscr{Y}$. Transition conditions are described by safety formulae of the form $\square \varphi_i$ where, for the environment $i = e$ and φ_e is a formula over $\mathscr{X} \cup \mathscr{X}' \cup \mathscr{Y}$, and for the system $i = s$ and φ_s is a formula over $\mathscr{X} \cup \mathscr{X}' \cup \mathscr{Y} \cup \mathscr{Y}'$. Note that the system plays second in each turn, so it can see $\mathscr{X}'$, whereas the environment cannot see $\mathscr{Y}'$, because it represents future values. The winning condition in a GR(1) game is described using progress formulae of the form $\square \diamond \psi_i, i \in \{e, s\}$, where ψ_i is a propositional formula over $\mathscr{X} \cup \mathscr{Y}$.

The overall specification of a GR(1) game is of the form

$$\underbrace{\left(\theta_e \wedge \square \varphi_e \wedge \bigwedge_{i=0}^{n-1} \square \diamond \psi_{e,i}\right)}_{assumption} \overset{\text{sr}}{\rightarrow} \underbrace{\left(\theta_s \wedge \square \varphi_s \wedge \bigwedge_{j=0}^{m-1} \square \diamond \psi_{s,j}\right)}_{assertion} \tag{1}$$

Note that requirements that constraint the environment are called *assumptions* and guarantees that the system must satisfy are called *assertions*. Assumptions limit the set of admissible environments, because, in practice, it is impossible to satisfy the design requirements in arbitrarily adversarial environments [6]. The *strict realizability* implication $\overset{\text{sr}}{\rightarrow}$ above is interpreted by prioritizing between safety and liveness [20], to prevent the system from violating the safety assertion, in case this would allow it to prevent the environment from satisfying the liveness assumption. The GR(1) synthesis algorithm has time complexity $O\left(nm|\Sigma|^2\right)$ [20], where n (m) is the number of recurrence assumptions (assertions), and Σ the set of all possible variable valuations.

3 Language definition

The language we are about to define is syntactically an extension of PROMELA [47], but its semantics is defined by a translation to turn-based infinite games with full information. PROMELA is a guarded command language that can represent transition systems, non-deterministic execution, and guard conditions for determining whether statements are executable [47, 27]. Its syntax can be found in the language reference manual [47, 46]. Here we will introduce syntactic elements only as needed for the presentation. Briefly, we mention that a program comprises of transition systems and automata, whose control flow

can be described with sequential composition, selection and iteration statements, `goto`, as well as blocks that group statements for atomic execution.

3.1 Variables

Ownership In a game, variables from $\mathscr{X}$ are controlled by the environment and variables from $\mathscr{Y}$ by the system. We call *owner* of a variable the player that controls it. We use the keywords `env` and `sys` to signify the owner of a variable. Variables can be of Boolean, bit, byte, (bounded) integer, or bitfield type.

Declarative and imperative semantics In imperative languages, variables remain unchanged, unless explicitly assigned new values. In declarative languages, variables are free to change, unless explicitly constrained [97]. In verification, both declarative languages like TLA [62] and SMV [24] have been used, as well as imperative languages like PROMELA and DVE [12].

In a synthesis problem, there are variables that are more succinct to describe declaratively, whereas others imperatively. For this reason, we combine the two paradigms, by introducing a new keyword `free` to distinguish between imperative and declarative variables. Variables whose declaration includes the keyword `free` are by default allowed to be assigned any value in their domain, unless explicitly constrained otherwise. Variables without the keyword `free` have imperative semantics, so their value remains unchanged, unless otherwise explicitly stated. Let V^{free} denote free variables, and V^{imp} imperative variables, and V_p the variables of player $p \in \{e, s\}$.

Ranged integer data type Symbolic methods for synthesis use reduced ordered binary decision diagrams (BDDs) [23, 10], which represent sets of states, and relations over states. As operations are performed between BDDs, these can grow quickly, consuming more memory. The growth can be ameliorated by reordering the variables over which a BDD is defined. Reordering variables can be prohibitively expensive, as discussed in Fig. 5, so reducing the number of bits is a primary objective. In addition, the complexity of GR(1) synthesis is polynomial in the number $|\Sigma|$ of variable valuations, which grows exponentially with each additional bit. We can reduce the number of bits by using bitfields whose width is tailored to the problem at hand. For convenience, the *ranged* integer type `int(MIN, MAX)` is introduced to define a variable $x \in \{\text{MIN}, \text{MIN}+1, \ldots, \text{MAX}\}$, with saturating semantics [42]. An integer with saturating semantics cannot be incremented when its value reached the maximal in its range, i.e., MAX.

A ranged integer is represented by a bitfield. The bitfield is comprised of bits, so it can only range between powers of two. The ranged integer though may have an arbitrary range. For this reason, safety constraints are automatically imposed on the bitfield representing the ranged integer. In other words, if x is the integer value of the bitfield, and it represents an integer that can take values from MIN to MAX, then the constraint $\square(\text{MIN} \leq x' \leq \text{MAX})$ is added to the safety formula, and $\text{MIN} \leq x \leq \text{MAX}$ to the initial condition of the player that owns the ranged integer.

Other numerical data types have mod wrap semantics. The value of an integer with mod wrap semantics overflows to MIN (underflows to MAX) if incremented when equal to the maximal value MAX (minimal value MIN). Mod wrap semantics are available only for integers that range over all values of a (signed) bitfield, because the modulo operation would otherwise be needed. Any BDD describing a modulo operation is at best of exponential size [23].

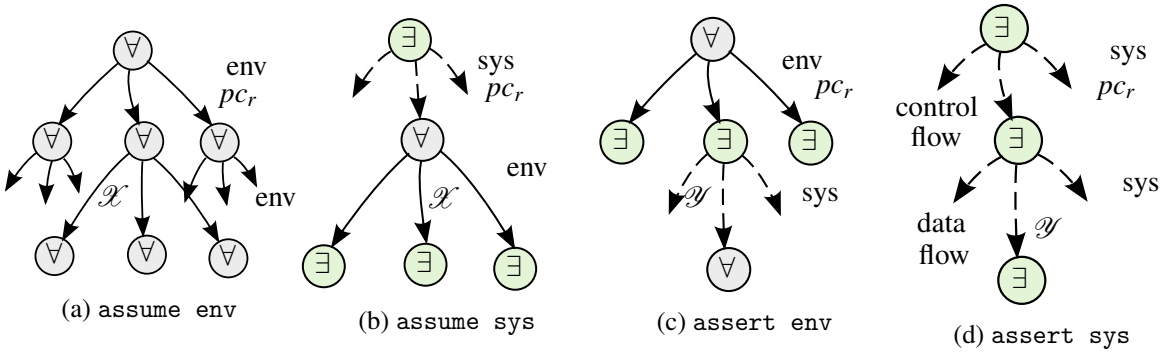

Figure 1: An assumption (assertion) process constrains the environment (system) variables, and `env` (`sys`) declares who chooses the next statement to be executed (when there are multiple).

3.2 Programs representing games

In many synthesis problems, the specification includes graph-like constraints. These may originate from physical configurations in robotics problems, deterministic automata to express a formula in GR(1), or describe abstractions of existing components that are to be controlled. These constraints can be described by processes. A process describes both control and data flow. In order to discuss control and data flow, we will refer to *program graphs*. A program graph is an intermediate representation of a process, after parsing and control flow analysis. For our purposes, a *program graph* is a rooted directed multi-graph $P_r \triangleq (V_r, E_r)$ whose edges E_r are labeled with program statements, and nodes V_r abstract states of the system [53, 10]. Execution starts from the graph's root. A multi-digraph is needed, because, between two given nodes, there may exist edges labeled with different program statements.

Control flow is the traversal of edges in a program graph (i.e., execution of statements), whereas data flow is the behavior of program variables along this traversal. A *program counter* pc_r is a variable used to store the current node in V_r. A natural question to ask is who controls the program counter. Another question is whose data flow is constrained by the program graph P_r. In the next section, we define syntax that allows declaring the player that controls the program counter, and the player that is constrained to manipulate the variables it owns, according to the statements selected by the program counter. This allows defining both processes where control and data flow are controlled and constrain the same player, but also processes with mixed control. If one player controls the program counter, and the opponent reacts by choosing a compliant data flow, then the process itself describes a game.

As an example, suppose that for a given process, the environment controls the program counter, and the system the local program variables. By choosing the next value of the program counter, the environment selects the next program statement that will execute. The system must react by choosing the next values of the local variables, such that they satisfy the selected program statement. The environment can select as next program statement any statement that is satisfiable by a system reaction, as discussed in more detail in Section 3.3. But other constraints, e.g., LTL formulae, can prevent the system from satisfying this statement.

The consecutive assignments of values to variables by the environment and the system can be represented by a *game graph*. Each node in the game graph corresponds to a valuation of variables. From each node, a single player can assign new values to its variables, depending on the outgoing edges at that node. The choice of outgoing edge at environment (system) nodes is universal (existential). The nodes in a game graph correspond to universal and existential nodes in alternating tree automata

[25, 79, 98, 99, 59]. Nondeterminism with universal quantification is known as *demonic*, [45, p.85], [95], otherwise as *angelic* [73, 37, 22]. In the previous example, the environment's choice of control flow occurs at universal nodes in the game graph. The system's reaction, by assigning to local variables, occurs at existential nodes in the game graph. The correspondence of control and data flow with a game graph is shown in Fig. 1.

3.2.1 Syntax

Program graphs are declared with the `proctype` keyword of `Promela` followed by statements enclosed in braces. The keyword `assume` (`assert`) declares a process that constrains the environment's (system's) data flow. These keywords are common in theorem proving and program verification languages [64].

The keyword `env` (`sys`) declares that the environment (system) controls the program counter pc_r of a process, Fig. 1. The implementation of `assume sys` is the most interesting, and is described in Section 4. We will call program graphs *processes*, noting that these processes have full information about each other, so they correspond to centralized synthesis, not distributed. The program counter *owner* is the player that controls variable pc_r. The *process player* is the player constrained by the program graph.

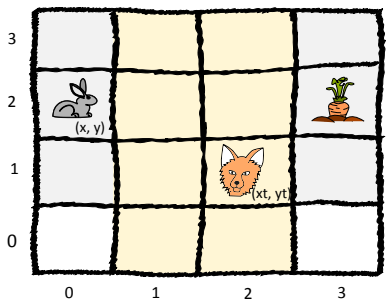

Figure 2: Adversarial game.

Example For example, the specification in Listing 1 defines a game between two players: the Bunny, and the Fox, that move in turns, as depicted in Fig. 2. Each logic time step includes a move by the Fox from (x_t, y_t) to (x_t', y_t'), followed by a move by the Bunny from (x, y) to (x', y'). The Bunny must reach the carrot, without moving through a cell that Fox is in (`assert ltl`). The Fox can only move between $x_t \in \{1, 2\}$, and has to keep visiting the lower row (`assume ltl`). The Fox can move diagonally, but the Bunny only vertically or horizontally. Both players have an option to stay still (`skip`). Note that x_t is a declarative variable, so it can change unless constrained.

Listing 1: Simple example.

```
1  #define H 3
2
3  free env int(1, 2) xt;
4  env int(0, H) yt;
5
6  assume env proctype fox(){
7      do
8      :: yt = yt - 1
9      :: yt = yt + 1
10     :: skip
11     od
12 }
13
14 assume ltl { []<>(yt == 0) }
15
16 sys int(0, 3) x;
17 sys int(0, H) y;
18
19 assert sys proctype bunny(){
20     do
21     :: x = x - 1
22     :: x = x + 1
23     :: y = y - 1
24     :: y = y + 1
25     :: skip
26     od
27 }
28
29 assert ltl {
30   [] ! ((x == xt) && (y == yt)
       ) &&
31   [] ! ((xt' == x) && (yt' ==
       y)) &&
32   []<>((x == 3) && (y == 2)) }
```

The process `fox` constrains the environment variables x_t, y_t (`assume`) and the environment controls its program counter (`env`). The `do` loops define alternatives that each player must choose from to continue playing the game. Note that nondeterminism in process `fox` is demonic (universally quantified), whereas in `bunny` angelic (existentially quantified), i.e., the design freedom given to the synthesis tool. Each player has full information about all variables in the game, both local, as well as global, and auxiliary. The solution is a strategy represented as a Mealy transducer [75] that the Bunny can use to win the game.

The conjuct $\varphi \triangleq \Box\neg((x = x'_t) \wedge (y = y'_t))$ prevents the Bunny from moving next to the Fox, from where the Fox can catch it in the next turn. The $\Box$ operator requires that, at each time step, the formula $\neg$ as main operator be true.

3.3 Statements

Control flow can be defined using selection (`if`) and repetition (`do`) statements, `else`, `break`, `goto`, and labeled statements. The statements `run`, `call`, `return` are not supported, because dynamic process creation would dynamically add BDD variables. In this section, we define expressions, assignments, and their executability.

Expressions Primed variables (that correspond to using the next operator $\bigcirc$) can appear in expressions to refer to the next values of those variables, as in the syntax of synthesis tools and TLA [61]. The operators weak previous $\circledcirc$ and strong previous $\ominus$ are expressed with the tokens `-X` and `--X`, respectively. Following TLA, we will call *(state) predicate* an expression that contains no primed variables and *action* an expression that contains primed variables [61]. Actions can be regarded as generalized assignments, in a sense that will be made precise later. Primed system variables cannot appear in assumption processes, because they refer to values not yet known to the environment. Using a GR(1) synthesizer as back-end, multiple priming within a single statement is not allowed, but can be allowed if a full LTL synthesizer is used as back-end [49, 36].

Deconstraining By default, imperative variables are constrained to remain invariant. If any assumption (assertion) process executes a statement that contains a primed environment (system) variable, then that variable is not constrained to remain unchanged in that time step. For example, in the assertion `sys bit x = 0; (x == 0); (x' == 1 - y)` the variable `x` is constrained by $x' = x$ when `x == 0` is executed, but the synthesizer is allowed to pick its next value as needed, in order to satisfy `x' == 1 - y`. Note that statements in assumption (assertion) processes that contain primed imperative system (environment) variables do *not* deconstrain those variables, because assumptions (assertions) are relevant only to the environment's (system's) data flow.

Assignments In PROMELA, expressions are evaluated by first converting all values to integers, then evaluating the expression with precision that depends on the operating system and processor, and updating the assigned variable's value, truncating if needed. Let $\text{trunc}(y, w)$ denote a function that truncates the value of expression y to bitwidth w. An assignment `x = expr` is translated to the logic formula $x' = \text{trunc}(expr, \text{width}(x))$, if variable x has mod wrap semantics, and to $x' = expr$ otherwise. If variable x is imperative, then it is deconstrained.

Statement executability A condition called *guard* is associated to each statement [27]. The process can execute a statement only if the guard evaluates to true. If a process currently has no executable

statement, then it *blocks*. For each statement, its guard is defined by existential quantification of the primed variables of the data flow player. The quantification is applied after the statement is translated to a logic formula. So the guard of a statement is the realizability condition for that statement. It means that, from the local viewpoint of that statement only, given the current values of variables in the game, the constrained player can choose a next move. So the scheduler cannot pick as next process to execute a process that has blocked. Clearly, if all processes block, then that player deadlocks.

Using this definition, the guard of a state predicate is itself, as in PROMELA. The implementation quantifies variables using the PYTHON binary decision diagram `dd` [4]. If an unsatisfiable guard is found, then the implementation raises a warning. For example, if we inserted the statement `xt && xt' && y'` in the process `bunny` (see example), then its guard would be $\exists y'.x_t \wedge x_t' \wedge y' = x_t \wedge x_t'$. Similarly, the guard of an expression `xt && xt'` in the process `fox` is x_t.

4 Translation to logic

In this section, we describe how a program is translated to temporal logic, in particular GR(1). For each process, the starting point is its program graph, which has edges labeled by program statements, and describes the control flow of a process in the source code. The construction of program graphs from source code is the same as for PROMELA [47], and described in detail in [35].

Here we give a brief example. Consider the process `maintain_lock` in Listing 2. It has two `do` loops, with two outgoing edges each. The corresponding program graph is shown in Fig. 3b. Each statement labels one edge, and that edge can be traversed if the guard associated to the statement evaluates to true. The guard can contain primed variables, requiring that the dataflow player manipulates them so as to make the edge's guard true. Otherwise, the player cannot traverse an edge with false guard. This program graph is translated further to logic, as described next. The semantics of the language are defined by this translation to logic.

There are three groups of elements in a program: processes, `ltl` blocks, and the scheduler that picks processes for execution. The scheduler is not present in the source code, but is added during translation, to represent the products between processes. The translation can be organized into a few thematically related sets of formulae. Due to lack of space, we are going to discuss the most interesting and representative of these at a high level. The full translation can be found in [35], and in the implementation. There are four groups of formulae: (i) control and data flow, (ii) invariance of variables, (iii) process scheduler, (iv) exclusive execution (atomic).

Control and data flow The translation of processes is reminiscent of symbolic model checking [74], but differs in that there are two players, and both play in each logic time step. This requires carefully separating the formulae into assumptions and guarantees (assertions).

Suppose that the scheduler selects process r to execute (how is explained later). At a given time step, a process is at some node i in its program graph, and will transition to a next node j, by traversing an edge labeled by a program statement. The player that controls the program counter pc_r selects the next statement, so the edge in the program graph. The player that is constrained by that process has to make sure that it *complies*, by picking values for variables that it controls such that the statement is satisfied. Recall that the scheduler can only pick from processes that have a satisfiable statement, so whenever a process executes, there will exist a satisfiable next statement. Of course, conflicts can arise between different synchronous processes that can lead to deadlock, and it is the synthesizer's task to avoid such

situations, to avoid losing the game. The transition constraints are encoded by the formula

$$\text{trans}(p,r) \triangleq \bigwedge_{i\in N_r} \big((pc_r = i) \rightarrow \bigvee_{(i,j,k)\in E_r} \varphi_{r,i,j,k} \wedge (\tilde{pc}_r = j) \wedge (\tilde{\text{key}}_r = k) \wedge \text{exclusive}(p,r,i,j,k)\ \big) \tag{2}$$

where N_r denotes the set of nodes, and E_r the multi-edges of the program graph of a process, p denotes the player (e,s). The logic formula equivalent to bitblasting the statement labeling edge (r,i,j,k) is $\varphi_{r,i,j,k}$. For `assume sys` processes, the system selects the next edge one time step before the scheduler decides whether that environment process will execute, so two copies are needed, $pc_r, \hat{pc}_r$ (system variables). So in an `assume sys` process, $\tilde{pc}_r \triangleq \hat{pc}_r, \tilde{\text{key}}_r \triangleq \text{key}(r)$, and in other processes $\tilde{pc}_r \triangleq pc'_r, \tilde{\text{key}}_r \triangleq \text{key}(r)'$. The variable $\text{key}(r)$ selects among multi-edges, and is controlled by the same player as the program counter pc_r of the process with pid r. In a system process, if node j is in an atomic block, then $\text{exclusive}(p,r,i,j,k)$ sets the auxiliary variables ex'_s and pm'_s to request atomic execution from the scheduler. The integer variable ex_s stores the identity of the process that requests atomic execution, and the bit pm_s requests that the environment be preempted, if the scheduler grants the request for atomic execution.

$$\begin{aligned}
\text{dataflow}(r) &\triangleq (ps(r)' = m(r)) \rightarrow \text{trans}(p,r)\\
\text{selectable}(r) &\triangleq \text{blocked}(r) \rightarrow (ps(r)' \neq m(r))\\
\text{control_flow}(r) &\triangleq \text{ite}\big((ps(r)' = m(r)),\ \text{pc_trans}(p,r),\ \text{inv}(pc_r)\big)\\
\text{blocked}(r) &\triangleq \bigvee_{i\in N_r} \big((pc_r = i) \wedge \bigwedge_{(i,j,k)\in E_r} \neg\text{guard}_{r,i,j,k}\big)
\end{aligned} \tag{3}$$

The environment variables $ps(r)$ select the process or synchronous product that will execute next inside an asynchronous product (top context is an asynchronous product). For this purpose, each process and product have a local integer id $m(\cdot)$ among the elements inside the product that contains them. The transition relation for the program counter depends on the type of process. For `assume sys` processes, $\text{pc_trans}(p,r) \triangleq (pc'_r = \widehat{pc}_r)$, and for other processes it equals $\text{guards}(r)$. The condition $\text{guards}(r)$ constrains the program counter to follow unblocked edges in a process. It is necessary when the control and data flow are controlled by different players, because whoever moves the program counter, can otherwise pick an edge with a statement that blocks the other player. In addition, for `assume sys` processes, a separate constraint with same form as $\text{guards}(r)$, but different priming of sub-expressions is imposed on the program counter copy $\widehat{pc}_r$. The ternary conditional is denoted by $\text{ite}(a,b,c)$.

Invariance of variables When a process is not executing, its declarative local variables must be constrained to remain invariant ($x' = x$). Also, imperative variables must remain invariant whenever no process executes a statement (edge) that either is an expression and contains a primed copy of that variable, or is an assignment. These are ensured by the following equations

$$\begin{aligned}
&\text{local_free}(p,r) \triangleq (ps(r)' \neq m(r)) \rightarrow \bigwedge_{x\in V^{\text{free}}_{p,r}} \text{inv}(x)\\
&\text{imperative_inv}(p,r) \triangleq \text{array_inv}(p,r) \wedge \bigwedge_{x\in V^{\text{imp}}_{p,r}} (\text{inv}(x) \vee \bigvee_{(i,j,k)\in E_r . x\in\text{deconstrained}(r,i,j,k)} \text{edge}(r,i,j,k))
\end{aligned} \tag{4}$$

where $V^{\text{free}}_{p,r}$ are free local variables of player p in process r. For `assume sys` processes, it is $\text{edge}(r,i,j,k) \triangleq (ps(r)' = m(r)) \wedge (pc_r = i) \wedge (\hat{pc}_r = j) \wedge (\text{key}(r) = k)$, and for other types of processes $\text{edge}(r,i,j,k) \triangleq (ps(r)' = m(r)) \wedge (pc_r = i) \wedge (pc'_r = j) \wedge (\text{key}(r)' = k)$. Primed references to elements in imperative arrays deconstrain only the referenced array element, ensured by array_inv.

Scheduler The scheduler (environment) has to select the processes that will execute. Products of processes can be defined in the source code by enclosing processes, or other products, in braces preceded by the keywords `async` and `sync`. They can be nested. `async` defines an asynchronous product, and the scheduler picks some unblocked process or product inside it to execute next. If all processes/products in an asynchronous product k have blocked, then the scheduler sets the corresponding variable ps_k to a reserved value (n_k). The reserved value is also used if the asynchronous product is nested in a synchronous product that currently is not selected to execute.

If the top product blocks, then the player has deadlocked, losing the game. The only exception is when the environment is preempted by a request from a system process for atomic execution. At a high level, this behavior is expressed as follows for scheduling the environment processes.

$$\begin{aligned}
\text{product_selected}(k) &\triangleq (ps(k)' \neq m(k)) \leftrightarrow (ps'_k = n_k) \\
\text{selectable_element}(r) &\triangleq \text{element_blocked(r)} \rightarrow (ps(r)' \neq m(r)) \\
\text{element_blocked}(r) &\triangleq \begin{cases} \text{blocked}(r), \text{ if } r \text{ is a process} \\ \text{sync_blocked}(r), \text{ if } r \text{ is a synchronous product} \\ \text{async_blocked}(r), \text{ if } r \text{ is an asynchronous product.} \end{cases}
\end{aligned} \tag{5}$$

$$\begin{aligned}
\text{sync_blocked}(k) &\triangleq \bigvee_{r \in R_k} \text{element_blocked}(r) \\
\text{async_blocked}(k) &\triangleq \bigwedge_{r \in R_k} \text{element_blocked}(r)
\end{aligned} \tag{6}$$

$$\text{pause_env_if_req} \triangleq (ps'_{\text{env_top}} = n_e) \leftrightarrow (pm_s \wedge (ps'_{\text{sys_top}} = ex_s < n_s)). \tag{7}$$

The expression element_blocked(r) depends on the blocked(z) expressions, and ensures that the scheduler doesn't select a synchronous product containing some blocked process, neither an asynchronous product where all processes are blocked. Recall also selectable from earlier, which applies to individual processes. Analogous formulae apply to system processes. For system processes, the top-level asynchronous product implication in async_blocked(r) must be replaced with equivalence, to force the environment to choose some system process (or product) to execute, when there exist unblocked ones. Note that the asynchronous products here are in the context of full information, so the system is *not* asynchronous in the sense of [55, 86].

Exclusive execution A system process of the top asynchronous product can request to execute atomically by setting the variables pm_s, ex_s, Eq. (2). If that process remains unblocked in the next time step, then the scheduler will grant it uninterrupted execution, until it exits atomic context (either blocked, or reached statements outside the `atomic{...}` block).

$$\text{grant}_s \triangleq \bigwedge_{r \in \text{pids}(s)} \big(((ex_s = m(r)) \wedge \text{frozen_unblocked}(r)) \rightarrow (ps(r)' = m(r))\big) \tag{8}$$

Recall also that the environment is allowed to pause only if preempted by the system, otherwise it loses the game (pause_env_if_req). The formula frozen_unblocked(r) checks whether the system would block, in case the environment froze, granting it exclusive execution. In case the system will block, then the request is not granted, and atomicity lost. This requires substituting primed environment variables with

unprimed ones, as follows

$$\begin{aligned}\text{frozen_unblocked}(r) &\triangleq \begin{cases} \bigvee_{i\in N_r} \big((pc_r = i) \wedge \bigvee_{(i,j,k)\in E_r} \text{guard_test}(r,i,j,k)\big), \text{ if player}(r) = s \\ \neg\text{blocked}(r), \text{ otherwise} \end{cases} \\ \text{guard_test}(r,i,j,k) &\triangleq \begin{cases} \text{guard}(r,i,j,k)|_{\,x/x' \text{ for } x\in\mathscr{X}}, \text{ if } i \text{ in atomic context} \\ \text{guard}(r,i,j,k), \text{ otherwise.} \end{cases}\end{aligned} \tag{9}$$

The reason is that this formula corresponds to the case that the environment sets $x' = x$ for the program variables it owns. If atomicity is lost in this turn, then the environment does not need to set $x' = x$, and this is ensured by the definition of guard_test(r,i,j,k).

As in PROMELA, LTL formulae that express safety are deactivated during atomic execution (in implementation, an option allows making atomic execution visible to LTL properties). They are re-activated as soon as atomicity is lost.

$$\text{mask_env_ltl} \triangleq \text{ite}(pm_s \wedge (ps'_{\text{sys_top}} = ex_s < n_s),\ \text{freeze_env_free},\ \psi_{\text{env safety } \texttt{ltl}}) \tag{10}$$

For the system, mask_sys_ltl is defined similarly. The formula freeze_env_free constrains declarative environment variables in global context and inside system processes to remain unchanged while the system is granted exclusive execution.

If an `atomic` block appears in a process, then the `ltl` properties in the program must not contain primed variables, to ensure that the above translation yields the intended interpretation (stutter invariance). If unbounded loops appear inside an atomic context, then there can be no liveness assumptions. The reason is that the system can hide in atomic execution forever, preventing the environment from satisfying its liveness assumptions, thus winning trivially. In order to avoid this, the environment liveness goals must be disjoined with strong fairness, a persistence property ($\Diamond\Box$), which is outside of the GR(1) fragment. An extension to use a full LTL synthesizer is possible, though not expected to scale as well. Labels in the code that contain progress result in accepting states (liveness conditions). Those expressions described but not defined above, the initial conditions, and a listing into assumptions and assertions can be found in the technical report [35].

5 Implementation

The implementation is written in PYTHON and available [1, 2, 5] under a BSD license. The frontend comprises of a parser generator that uses PLY (Python `lex-yacc`) [14]. The parser for the proposed language subclasses and extends a separate parser for PROMELA [2], to enable use of the latter also by those interested in verification activities. After parsing and program graph construction, the translation described in Section 4 is applied [1]. This results in linear temporal logic formulae that contain modular integer arithmetic. At this point, each `ltl` block is syntactically checked to be in the GR(1) fragment, and split into initial condition, action, and recurrence conjuncts [5]. The past fragment is then translated using temporal testers [54]. In the future, the syntactic check can be removed, and a full LTL synthesis algorithm used.

The next step encodes signed arithmetic in bitvector logic using two's complement representation [58]. The resulting formulae are in the input syntax recognized by the SLUGS synthesizer [32]. This prefix syntax includes *memory buffers*, which enable avoiding repetition of formulae. For example, `$ 3 x a b & ?1 ?0 | ?2 ! ?0` describes the ternary conditional ite(x,a,b). Memory buffers

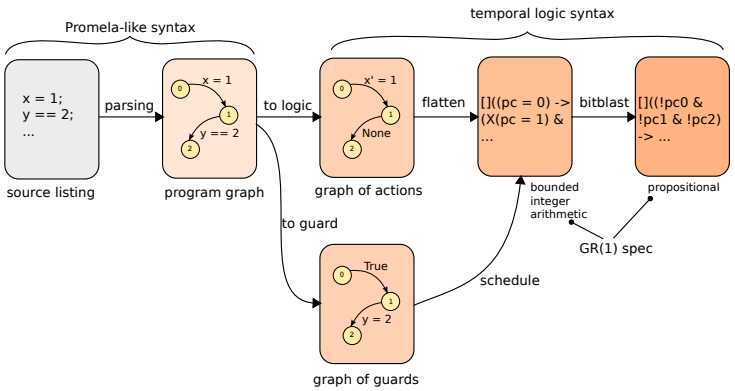

(a) Compiler architecture.

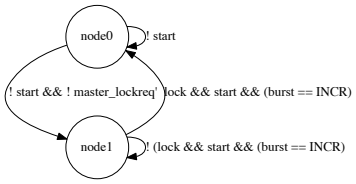

(b) Program graph that corresponds to process `maintain_lock` in Listing 2.

Figure 3: Compiling programs to temporal logic.

prevent the bitblasted formulae from blowing up. The SLUGS distribution includes an encoder of unsigned addition and comparison into bitvector logic using memory buffers. Here, signed arithmetic and arrays are supported. The bitblaster code is a separate module, which can be reused as a backend to other frontends. The resulting formula is passed to the SLUGS synthesis tool to check for realizability and construct a winning strategy as a Mealy transducer.

6 AMBA AHB Case study

Revised specification The ARM processor Advanced Microcontroller Bus Architecture (AMBA) [9] specifies a number of different bus protocols. Among them, the Advanced High-performance (AHB) architecture has been studied extensively in the reactive synthesis literature [17, 18, 77, 20, 91, 43, 19].

The AHB bus comprises of masters that need to communicate with slaves, and an arbiter that controls the bus and decides which master is given access to the bus. The arbiter receives requests from the masters that desire to access the bus, and must respond in a weakly fair way. In other words, every master that keeps uninterruptedly requesting the bus must eventually be granted access to it. Note that the AMBA technical manual [9] does not specify any fairness requirement, but instead leaves that decision to the designer. For automated synthesis, weak fairness is one possible formalization that ensures servicing of all the masters.

In addition, a master can request that the access be locked. In the ARM manual, the arbiter makes no promises as to whether a request for the lock will be granted. If the arbiter does lock the access, then it guarantees to maintain the lock, until the request for locking is withdrawn by the master that currently owns the bus. Note that the specification used here requires the arbiter to lock the bus, whenever requested by the master to be granted next.

A specification for the arbiter appeared in [17], and is presented in detail in [20]. Here, we expressed the specification of [20] in the proposed language, Listing 2 on p.90. In doing so, some assumptions were weakened and assumption A1 modified, to improve the correspondence with the ARM technical manual, and reduce the number of environment variables (thus universal branching). First, we describe the AHB specification, referring to Listing 2. After that, we summarize the changes, and discuss the experiments.

The specification in Listing 2 has both environment and system variables, as well as assumptions and guarantees. The arbiter is the system, and the environment comprises of slaves and $N+1$ masters. An

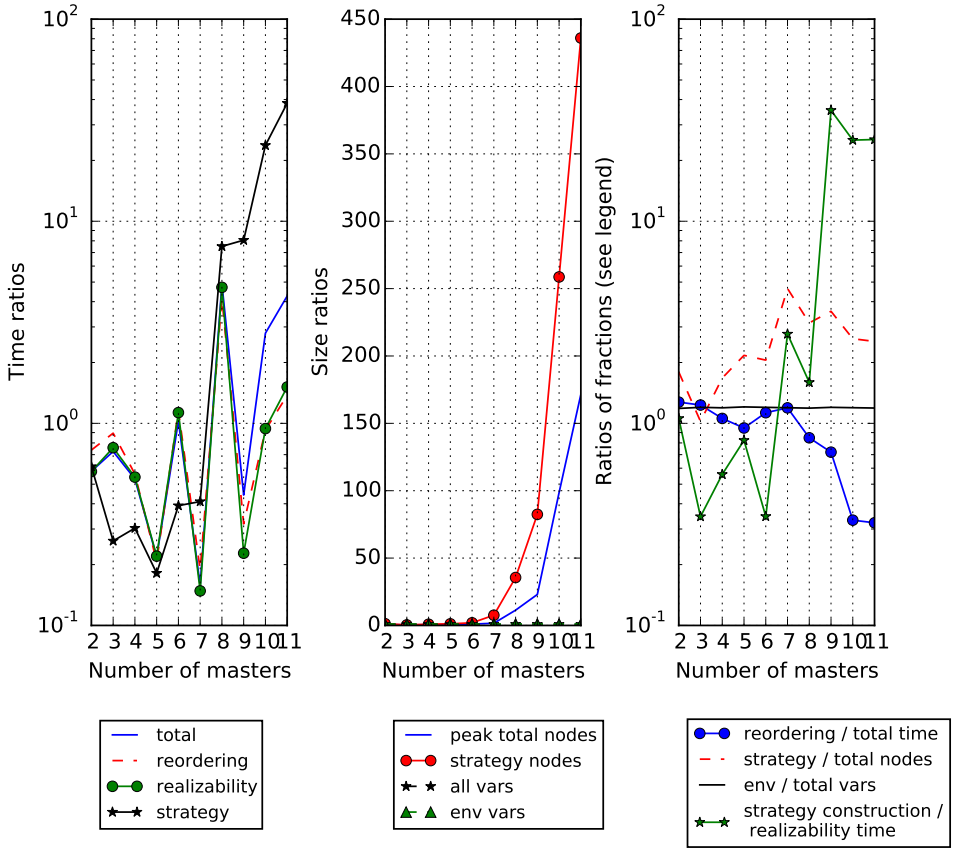

(a) Conjunction divided by BA, no reordering.

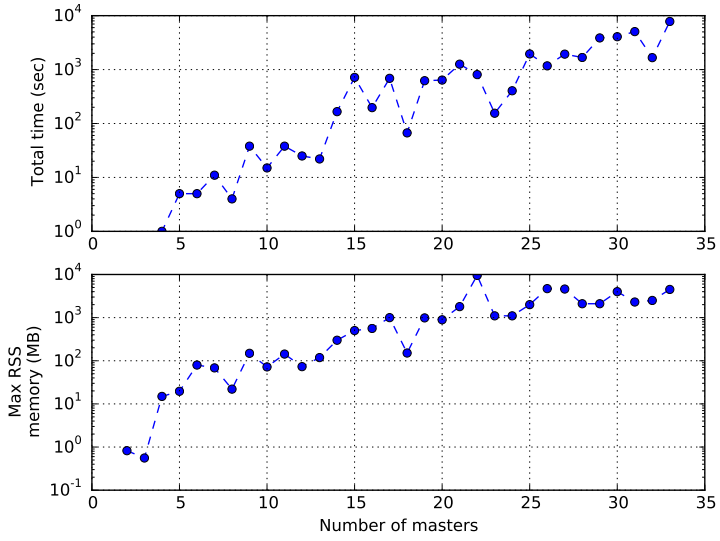

(b) BA, no reordering.

Figure 4: Selection of experimental measurements for the revised AMBA specification.

array `request` of bits represents the request of each individual master to be given bus ownership, for sending and receiving from a slave of interest. Communication proceeds in bursts. The bus owner selects which type of burst it desires, by setting the integer `burst`. Three lengths of bursts are modeled: single time step (`SINGLE`), four consecutive time steps (`BURST4`), and undefined duration (`INCR`). The currently addressed slave sets the bit `ready` (to true) to acknowledge that it has successfully received data for a burst. While `ready` is false, the bus owner cannot change (G1,6), and a `BURST4` time step is not counted towards completion (G3). For this reason, the slaves (environment) are required to recur setting `ready` to true (A2). The master can also request that, when it is granted the bus, it should be locked. Each master can do so by setting a signal. Only two of these signals are modeled here, using the bit variables `grantee_lockreq` and `master_lockreq` (described below).

The arbiter works in primarily two phases, as introduced in [20]. These phases are extraneous to the standard, and used only to aid in describing the specification. Firstly, the arbiter decides to which master it will next grant the bus to. The arbiter sets the bit `decide` to true during that period. The decision is stored in the form of two variables, `grant` and `lockmemo`, which don't change while `decide` is false (G8). The integer `grant` indicates the master that has been decided to receive bus ownership after the current owner. For performance reasons, the arbiter can only grant the bus to a master that requested it (G10), with the exception of a default master (with index 0).

The bit `lockmemo` is set to the value of the environment bit `grantee_lockreq` (G7). The value `grantee_lockreq'` represents whether the master `grant` had requested locked ownership. In the original specification, an array *lockreq* of *N* environment bits is used (denoted by *HLOCK* in [20]). This increases significantly the variables with universal quantification. Here, this array is abstracted by the bit `grantee_lockreq`. In implementation, the transducer input `grantee_lockreq'` should be set equal to the lock request of master `grant` in the previous time step, i.e., $grantee_lockreq' \triangleq (\ominus lockreq)[grant]$. In [20], some assumptions are expressed to constrain the array *lockreq*, i.e., when masters request locked ownership. The assumptions can be weakened [34], and by modifying assumption A1 (described below), the array *lockreq* can be abstracted by the two bits `grantee_lockreq` and `master_lockeq`.

The arbiter promises to lock the bus, until the bus owner `master` interrupts requesting it. The owner indicates its lock request by the value `master_lockreq'`. In implementation, the input value `master_lockreq'` should be set equal to the lock request of `master`, i.e., $master_lockreq' \triangleq lockreq[master]$.

In the second phase, the master changes the bus owner, by updating the integer `master` to `grant` (G4,5). If the grantee had requested a lock, via `grantee_lockreq`, then that request is propagated to the bit `lock` (G4,5). With the bit `lock`, the arbiter indicates that `master` has been given locked access to the bus.

To be weakly fair, each master that keeps uninterruptedly requesting the bus should be granted ownership. This requirement is described as a Büchi automaton (G9).

The assumption A1 of [20] requires that for locked undefined-length bursts, the masters eventually withdraw their request to access the bus. This assumption is not explicit in the ARM standard, so we modify it, by requiring that masters withdraw only their request for the lock, *not* for bus access. This is described as the Büchi automaton `withdraw_lock` that constrains the environment. The arbiter grants `master` locked access by setting the bit `lock` to true. If `lock` is false, then the master (environment) remains in the outer loop, at the `else`. If `lock` becomes true, then the automaton enters the inner loop. In order for the automaton `withdraw_lock` to exit the inner loop, the environment must set `master_lockreq'` to false. This obliges the owner `master` to eventually stop requesting locked ownership.

For a `SINGLE` burst, the burst is completed at the next time step that `ready` is true, so the arbiter does not need to lock the bus (since the owner remains unchanged while `ready` is false). For a `BURST4`

burst, the arbiter locks the bus for a predefined length of four successful beats (G3). This requirement is described by the safety automaton `count_bursts`. Note that assumption A1 is not needed for this case. For a `INCR` burst, the duration is unspecified a priori. While the owner `master` continuously requests locked access (with `master_lockreq`), the arbiter cannot change the bus owner (G2). This is described by the safety automaton `maintain_lock`. When the arbiter grants locked access to the bus for a burst of undefined duration, then the guard `lock && start && (burst == INCR)` is true. The process `maintain_lock` enters the inner loop, and remains there until `master_lockreq` becomes true. This is where assumption A1 is required, to ensure that the owner will eventually stop requesting the lock. The arbiter can then exit the inner loop of `maintain_lock`. Then, the arbiter can wait outside (`start` is false throughout the burst), until the addressed slave sets `ready` to true, signifying the successful completion of that burst, and allowing the arbiter to set `start` and change the bus owner, if needed.

Some properties not in GR(1) are translated to deterministic Büchi automata in [20]. The resulting formulae are much less readable, and not easy to modify and experiment with. Above, we specified these properties directly as processes, with progress states where needed.

Observations By encoding `master` and `grant` as integers, and abstracting the array *lockreq* by the two variables `master_lockreq` and `grantee_lockreq`, the synthesis time was reduced significantly (by a factor of 100 [34]), but are not sufficient to prevent the synthesized strategies from blowing up. By also merging the N weak fairness guarantees $\bigwedge_{i=0}^{N-1} \square\diamondsuit(request[i] \rightarrow master = i)$ into the Büchi automaton (BA) `weak_fairness` with one accepting state, we were able to prevent the strategies from blowing up, and synthesize up to 33 masters, Fig. 4b. The synthesis time for 16 masters is in the order of 5 minutes, and peak memory consumption less than 1GB. To our knowledge, in previous works, the maximal number of masters has been 16, the strategies were blowing up, and the runtimes were significantly longer (21 hours for 12 masters in [20], and more than an hour in [43] for 16 masters).

Measurements To identify what caused this difference, we conducted experiments for 8 different combinations: original vs revised spec, conjunction vs BA, reordering during strategy construction enabled/disabled, Table 1. We collected detailed measurements with instrumentation that we inserted into SLUGS, available at [3]. Some of these are shown in Fig. 5, and the complete set can be found in the technical report [34] (the language is described in [35]). The experiments were run on an Intel(R) Xeon® X5550 core, with 27 GB RAM, running Ubuntu 14.04.1. The maximal memory limit of CUDD [94] was set to 16 GB.

We found that lack of dynamic BDD reordering during construction of the strategy was the reason for poor performance of conjoined liveness goals, as opposed to a single BA. The implementation of the GR(1) synthesis algorithm in SLUGS has three phases:

1. Computing the winning region, while memoizing the iterates of the fixpoint iteration, as BDDs.
2. Construction of individual strategies, one for each recurrence goal.
3. Combination of the individual strategies into a single transducer, which iterates through them.

In SLUGS, variable reordering [90] is enabled during the first two phases, but disabled in the last one. If the recurrence goals are conjoined into a formula of the form $\bigwedge\square\diamondsuit$, then the memory needed for synthesis blows up Fig. 4a, for both the original and revised specifications. Using a BA, the revised specification scales without blowup.

If reordering is enabled during the last phase (combined transducer construction), then the specification with conjunction can be synthesized without blowup. With a BA, enabling reordering in the last

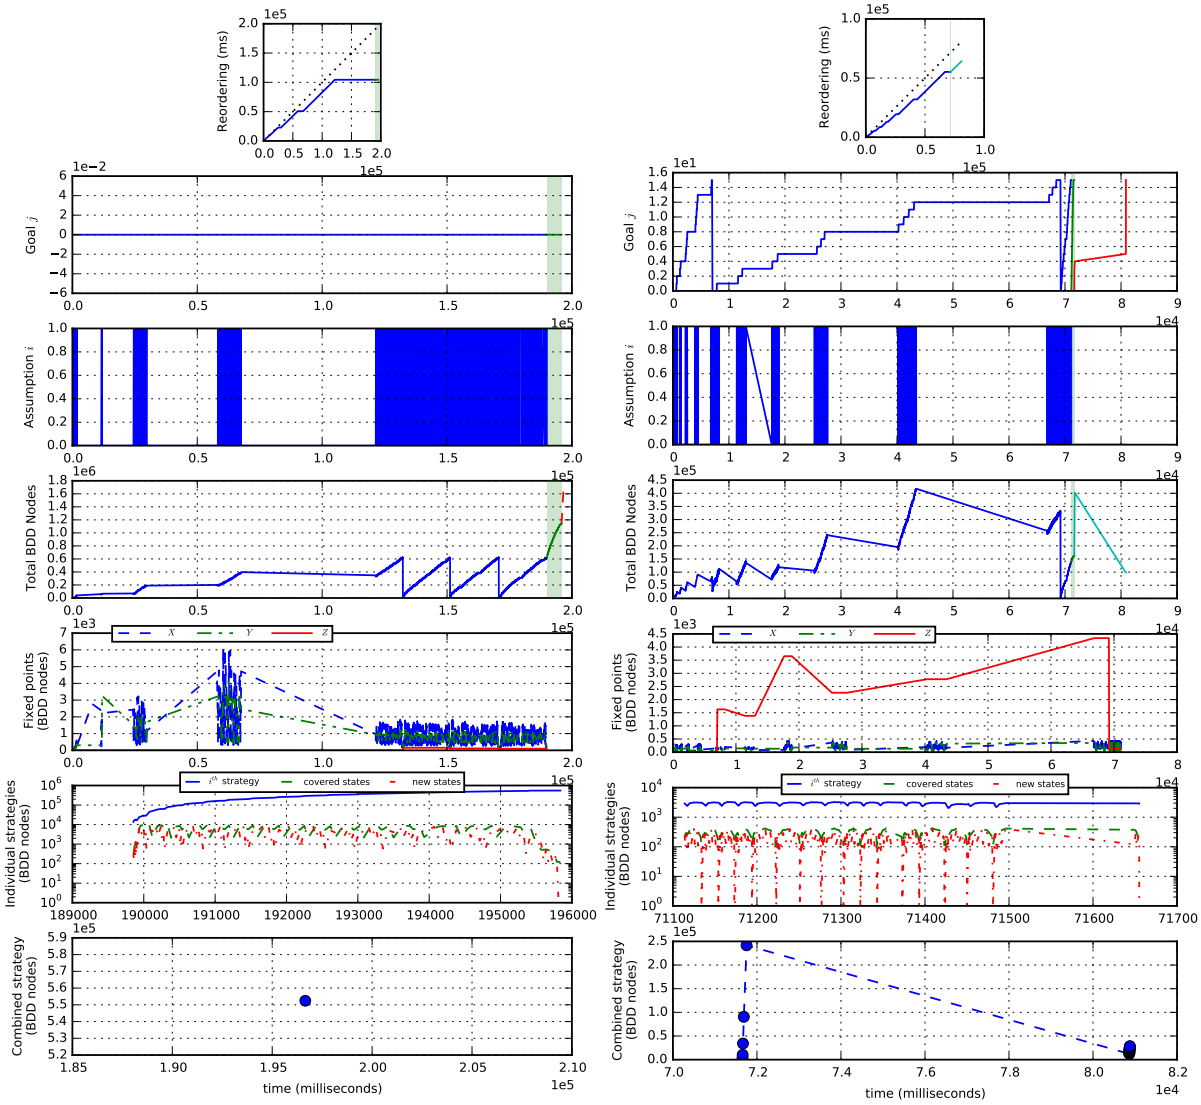

(a) BA, no reordering, $N = 16$.

(b) Conjunction, with reordering, $N = 16$.

Figure 5: Measurements during phases of: (1) realizability, (2) sub-strategy, and (3) combined strategy construction. The top 4 plots are over all phases, the fixpoints over (1), the individual strategies over (2), and the bottom plot over (3). The revised specification is used.

phase has a mildly negative effect, because it can trigger unnecessary reordering. We used the group sifting algorithm [82, 90] for reordering.

Enabling dynamic BDD variable reordering is necessary to prevent the blowup. The conjunction with reordering enabled in phase 3 outperforms the BA with reordering turned off in phase 3. This is a consequence mainly of the fact that the BA chains the goals inside the state space, leading to deeper fixpoint iterations, and has slightly larger state space, due to the nodes of the automaton `maintain_lock`.

Reordering typically accounts for most of the runtime (top plot in Fig. 5a). The second plot shows the currently pursued goal during realizability, and later the sub-strategy under construction, and the sub-strategy being combined in the final strategy. Each drop in total BDD nodes (teeth in 4th plot)

Table 1: Overview of results.

	Strategy reordering	Specification original	Specification revised
Conjunction of fairness	with	slow	fast
	w/o	memory blowup	memory blowup
Büchi automaton	with	very slow	ok (slower)
	w/o	slow	ok

corresponds to an outer fixpoint iteration. The first outer iteration takes the most time, due to reordering. Later iterations construct subsets, for which the obtained order remains suitable. The highlighted period corresponds to the construction of individual strategies. Plots for the other experiments can be found in [34].

Conclusions The major effect of reordering in the final phase of strategy construction can be understood as follows. Using a BA reduces the goals to only one, so no disjunction of individual sub-strategies is needed [83]. Also, this encoding shifts the transducer memory (a counter of liveness goals), from the strategy construction, to the realizability phase (attractor computations). This slightly increases the state space. Nonetheless, this symbolic encoding allows the variable ordering more time to gradually adjust to the represented sets.

In contrast, by conjoining liveness goals, the variable order is oblivious during realizability checking that the sub-strategies will be disjoined at the end. The disjunction of strategies acts as a shock wave, disruptive to how far from optimal the obtained variable order is. If, by that phase, reordering has been disabled, then this effect causes exponential blowup.

Overall, the proposed language made experimentation easier and revisions faster, helping to study variants of the specification. It can be used to explore the sensitivity of a specification, in the following way. A formula, e.g., requiring weak fairness, can be temporarily replaced with a process that is one possible refinement of that formula, potentially simplified. In the AMBA example, one can fix a round robin schedule for selecting the next grantee (temporarily dropping G10). This is reminiscent of the manual implementation [20]. By doing so, it can be evaluated whether the synthesizer finds it difficult to pick requestors only, or whether some other factor is more important, either another part of the specification, or some external factor. For the AMBA problem, such a change resulted in replacing recurrence formulae with a BA, and led to identifying the need for strategy reordering to avoid memory blowup. Therefore, we believe that it can prove useful in exploring the sensitivity of specifications, to help the specifier direct their attention to improve those parts of the specification that impact the most synthesis performance.

Listing 2: AMBA AHB specification in the proposed language.

```
1  #define N 2  /* N + 1 masters
       */
2  #define SINGLE 0
3  #define BURST4 1
4  #define INCR 2
5  /* variables of masters and
       slaves
6  A4: initial condition */
7  free env bool ready = false;
8  free env int(0, 2) burst;
9  free env bool request[N + 1] =
       false;
10 free env bool grantee_lockreq=
```

```
    false;
free env bool master_lockreq =
    false;
/* arbiter variables */
/* G11: sys initial condition
    */
free bool start = true;
free bool decide = true;
free bool lock = false;
free bool lockmemo;
free int(0, N) master = 0;
free int(0, N) grant;
/* A2: slaves must progress
    with receiving data */
assume ltl { []<> ready }
/* A3: dropped, weakening the
    assumptions */
/* A1: */
assume env proctype
    withdraw_lock(){
     progress:
     do
     :: lock;
         do
         :: ! master_lockreq';
            break
         :: true /* wait */
         od
     :: else
     od
}
assert ltl {
[](
/* G1: new access starts only
    when slave is ready */
(start' -> ready)
/* G4,5 */
&& (ready -> ((master' ==
    grant) && (lock' <->
    lockmemo')))
/* G6 */
&& (! start' -> (
     (master' == master) &&
     (lock' <-> lock)))
/* G7: remember if lock
    requested */
&& ((--X decide) -> (lockmemo'
     <-> grantee_lockreq'))
/* G8 */
&& (! decide -> (grant' ==
    grant))
&& ((! --X decide) -> (
    lockmemo' <-> lockmemo))
/* G10: grant only to
    requestors */
&& ((grant' == grant) || (
    grant' == 0) || request[
    grant'])
)
}
sync{ /* synchronous product
    */
/* G9: weak fairness */
assert proctype weak_fairness
    (){
     int(0, N) count;
     do
     :: (! request[count] || (
        master == count));
         if
         :: (count < N) && (
            count' == count +
            1)
         :: (count == N) && (
            count' == 0);
             progress: skip
         fi
     :: else
     od
}
/* G2: lock until no lock req
    */
assert sys proctype
    maintain_lock(){
     do
     :: (lock && start && (
        burst == INCR));
         do
         :: (! start && !
            master_lockreq');
            break
         :: ! start
         od
     :: else
     od
}
/* G3: for a BURST4 access,
    count the "ready" time
    steps. */
assert sys proctype
    count_burst(){
     int(0, 3) count;
     do
     :: (start && lock &&
```

```
84          (burst == BURST4) &&
85          (!ready || (count' ==
               1)) &&
86          (ready || (count' ==
               0)) );
87          do
88          :: (! start && ! ready
               )
89          :: (! start && ready
               && (count < 3) &&
90               (count' == count +
                    1))
91          :: (! start && ready
               && (count >= 3));
               break
92          od
93      :: else
94      od
95  }
96  }
```

7 Relevant work

Our approach has common elements with program repair [51], program sketching [65], and syntax-guided synthesis [7]. Program repair aims at modifying an existing program in a conventional programming language. Syntax-guided synthesis uses a grammar to slice the admissible search space of terminating programs. Here, we are interested in reactive programs. Similarly, program sketching uses templates to restrict the search space and give hints to the synthesizer for obtaining a complete program. In [15], the authors propose another constraint-based approach to games, but start directly from logic formulae.

TLA [61, 62] subsumes our proposed language, since it includes quantification, but is intended as a theorem proving activity, is declarative, and is aimed at verification. Nonetheless, one can view the proposed translation as from open-PROMELA to TLA. SMV is a declarative language [24], and JTLV [88] an SMV-like language for synthesis specifications, but with no imperative constructs. ASPECTLTL is a further declarative extension for aspect-oriented programming [72].

RPROMELA is an extension of PROMELA that adds synchronous-reactive constructs (not in the sense of reactive synthesis) that include synchronous products and channels called ports [80, 81]. Its semantics are defined in terms of *stable states*, where the synchronous product blocks, waiting for message reception from its global ports. RPROMELA does not address modeling of the environment, nor declarative elements. Besides, synchronous-reactive languages like ESTEREL, QUARTZ (imperative textual), STATECHARTS, ARGOS, SYNCCHARTS (imperative graphical), LUSTRE, and LUCID SYNCHRONE (declarative textual) and SIGNAL (declarative graphical) are by definition *deterministic* languages intended for direct design of transducers [41, 44, 52]. In synthesis, non-determinism is an essential feature of the specification.

Our approach has common elements with *constraint imperative programming* (CIP), introduced with the experimental language KALEIDOSCOPE [38, 39, 40, 68], one of the first attempts to integrate the imperative and declarative constraint programming paradigms. An observation from [38], which applies also here, is that specifiers need to express two types of relations: long-lived (best described declaratively), and sequencing relations (more naturally expressed in an imperative style). However, CIP does *not* ensure correct reactivity, because the constraints are solved online. Constraints are a related approach that uses constraints for indirect assignment to imperative variables is [63].

The translation from PROMELA to declarative formalisms has been considered in [11, 48, 26] and decision diagrams in [13]. These translations aim at verification, do not have LTL as target language, and either have limited support for atomicity [26], no details [11], or programs graphs semantics that do not match PROMELA [48].

8 Conclusions

We have presented a language for reactive synthesis that combines declarative and imperative elements to allow using the most suitable paradigm for each requirement, to write readable specifications. By expressing the AMBA specification in a multi-paradigm language, it became easier to experiment and transform it into one that led to efficient synthesis that improved previous results by two orders of magnitude. Besides the AMBA specification, other examples can be found in the code repository [1].

Acknowledgments The authors would like to thank Scott Livingston for providing helpful feedback. This work was supported by STARnet, a Semiconductor Research Corporation program, sponsored by MARCO and DARPA. The first author was partially supported by a graduate research fellowship from the Jet Propulsion Laboratory, over the summer of 2014.

References

[1] *open*-PROMELA *compiler (*PYTHON *package)*. Available at `https://github.com/johnyf/openpromela`.

[2] PROMELA *parser (*PYTHON *package)*. Available at `https://github.com/johnyf/promela`.

[3] SLUGS *instrumentation at tag* `synt_2015`. Available at `https://github.com/johnyf/slugs`.

[4] `dd`*: Decision diagrams (*PYTHON *package)*. Available at `https://github.com/johnyf/dd`.

[5] `omega`*: Symbolic and enumerated data structures and algorithms for manipulating* ω*-regular sets (*PYTHON *package)*. Available at `https://github.com/johnyf/omega`.

[6] Martín Abadi & Leslie Lamport (1994): *Open systems in TLA*. In: *PODC*, pp. 81–90, doi:10.1145/197917.197960.

[7] Rajeev Alur, Rastislav Bodik, Garvit Juniwal, Milo MK Martin, Mukund Raghothaman, Sanjit A Seshia, Rishabh Singh, Armando Solar-Lezama, Emina Torlak & Abhishek Udupa (2013): *Syntax-guided synthesis*. In: *FMCAD*, pp. 1–17, doi:10.1109/FMCAD.2013.6679385.

[8] Rajeev Alur & Salvatore La Torre (2004): *Deterministic generators and games for LTL fragments*. *ACM Trans. Comput. Logic* 5(1), pp. 1–25, doi:10.1145/963927.963928.

[9] ARM Ltd. (1999): *AMBA™ Specification*, Rev 2.0 edition. Available at `http://www-micro.deis.unibo.it/~magagni/amba99.pdf`.

[10] Christel Baier & Joost-Pieter Katoen (2008): *Principles of model checking*. The MIT Press.

[11] Michael Baldamus & Jochen Schröder-Babo (2001): *P2B: A translation utility for linking* PROMELA *and symbolic model checking*. In: *SPIN*, pp. 183–191, doi:10.1007/3-540-45139-0_11.

[12] Jiří Barnat, Luboš Brim, Vojtěch Havel, Jan Havlíček, Jan Kriho, Milan Lenčo, Petr Ročkai, Vladimír Štill & Jiří Weiser (2013): DIVINE *3.0 – An explicit-state model checker for multithreaded C & C++ programs*. In: *CAV*, 8044, pp. 863–868, doi:10.1007/978-3-642-39799-8_60.

[13] Vincent Beaudenon, Emmanuelle Encrenaz & Sami Taktak (2010): *Data decision diagrams for promela systems analysis*. *STTT* 12(5), pp. 337–352, doi:10.1007/s10009-010-0135-0.

[14] David M. Beazley: *PLY (Python Lex-Yacc) v3.4*. Available at `http://www.dabeaz.com/ply/ply.html`.

[15] Tewodros Beyene, Swarat Chaudhuri, Corneliu Popeea & Andrey Rybalchenko (2014): *A constraint-based approach to solving games on infinite graphs*. In: *POPL*, pp. 221–233, doi:10.1145/2535838.2535860.

[16] Roderick Bloem, Alessandro Cimatti, Karin Greimel, Georg Hofferek, Robert Könighofer, Marco Roveri, Viktor Schuppan & Richard Seeber (2010): RATSY–*A new requirements analysis tool with synthesis*. In: *CAV*, pp. 425–429, doi:10.1007/978-3-642-14295-6_37.

[17] Roderick Bloem, Stefan Galler, Barbara Jobstmann, Nir Piterman, Amir Pnueli & Martin Weiglhofer (2007): *Interactive presentation: Automatic hardware synthesis from specifications: A case study*. In: *Design, Automation and Test in Europe (DATE)*, pp. 1188–1193. Available at `http://dl.acm.org/citation.cfm?id=1266366.1266622`.

[18] Roderick Bloem, Stefan Galler, Barbara Jobstmann, Nir Piterman, Amir Pnueli & Martin Weiglhofer (2007): *Specify, compile, run: Hardware from* PSL. *ENTCS* 190(4), pp. 3–16, doi:10.1016/j.entcs.2007.09.004.

[19] Roderick Bloem, Swen Jacobs & Ayrat Khalimov (2014): *Parameterized synthesis case study: AMBA AHB*. In Krishnendu Chatterjee, Rüdiger Ehlers & Susmit Jha, editors: *SYNT*, *EPTCS* 157, pp. 68–83, doi:10.4204/EPTCS.157.9. Available at `http://arxiv.org/abs/1407.6580v1`.

[20] Roderick Bloem, Barbara Jobstmann, Nir Piterman, Amir Pnueli & Yaniv Sa'ar (2012): *Synthesis of reactive(1) designs*. *Journal of Computer and System Sciences (JCSS)* 78(3), pp. 911–938, doi:10.1016/j.jcss.2011.08.007.

[21] Aaron Bohy, Véronique Bruyère, Emmanuel Filiot, Naiyong Jin & Jean-François Raskin (2012): ACACIA+, *a tool for LTL synthesis*. In: *CAV*, pp. 652–657, doi:10.1007/978-3-642-31424-7_45.

[22] Manfred Broy (1986): *A theory for nondeterminism, parallelism, communication, and concurrency*. *TCS* 45(0), pp. 1–61, doi:10.1016/0304-3975(86)90040-X.

[23] Randal E. Bryant (1986): *Graph-based algorithms for Boolean function manipulation*. *IEEE Trans. Comput.* 35(8), pp. 677–691, doi:10.1109/TC.1986.1676819.

[24] Roberto Cavada, Alessandro Cimatti, Charles Arthur Jochim, Gavin Keighren, Emanuele Olivetti, Marco Pistore, Marco Roveri & Andrei Tchaltsev (2010): NUSMV *2.5 User Manual*. Technical Report, Fondazione Bruno Kessler, 18 Via Sommarive, 38055 Povo (Trento), Italy.

[25] Ashok K. Chandra, Dexter C. Kozen & Larry J. Stockmeyer (1981): *Alternation*. *JACM* 28(1), pp. 114–133, doi:10.1145/322234.322243.

[26] Frank Ciesinski, Christel Baier, Marcus Größer & David Parker (2008): *Generating compact MTBDD-representations from Probmela specifications*. In: *SPIN*, pp. 60–76, doi:10.1007/978-3-540-85114-1_7.

[27] Edsger W. Dijkstra (1975): *Guarded commands, nondeterminacy and formal derivation of programs*. *CACM* 18(8), pp. 453–457, doi:10.1145/360933.360975.

[28] M.B. Dwyer, G.S. Avrunin & J.C. Corbett (1999): *Patterns in property specifications for finite-state verification*. In: *ICSE*, pp. 411–420, doi:10.1145/302405.302672.

[29] Rüdiger Ehlers (2011): *Experimental aspects of synthesis*. *EPTCS* 50, doi:10.4204/EPTCS.50.

[30] Rüdiger Ehlers (2011): *Generalized Rabin(1) synthesis with applications to robust system synthesis*. In: *NFM*, pp. 101–115, doi:10.1007/978-3-642-20398-5_9.

[31] Rüdiger Ehlers (2011): *Unbeast: Symbolic bounded synthesis*. In: *TACAS*, pp. 272–275, doi:10.1007/978-3-642-19835-9_25.

[32] Rüdiger Ehlers & Vasumathi Raman (2014): *Low-effort specification debugging and analysis*. *EPTCS* 157, pp. 117–133, doi:10.4204/EPTCS.157.12.

[33] Ioannis Filippidis & contributors (2013): *List of verification and synthesis tools*. Available at `https://github.com/johnyf/tool_lists/blob/master/verification_synthesis.md`.

[34] Ioannis Filippidis & Richard M. Murray (2015): *Revisiting the AMBA AHB bus case study*. Technical Report CaltechCDSTR:2015.004, California Institute of Technology, Pasadena, CA. Available at `http://resolver.caltech.edu/CaltechCDSTR:2015.004`.

[35] Ioannis Filippidis, Richard M. Murray & Gerard J. Holzmann (2015): *Synthesis from multi-paradigm specifications*. Technical Report CaltechCDSTR:2015.003, California Institute of Technology, Pasadena, CA. Available at `http://resolver.caltech.edu/CaltechCDSTR:2015.003`.

[36] Bernd Finkbeiner & Sven Schewe (2013): *Bounded synthesis*. *International Journal on Software Tools for Technology Transfer (STTT)* 15(5-6), pp. 519–539, doi:10.1007/s10009-012-0228-z.

[37] Robert W. Floyd (1967): *Nondeterministic algorithms*. *JACM* 14(4), pp. 636–644, doi:10.1145/321420.321422.

[38] Bjorn N. Freeman-Benson (1990): KALEIDOSCOPE*: Mixing objects, constraints, and imperative programming*. In: *OOPSLA/ECOOP*, pp. 77–88, doi:10.1145/97946.97957.

[39] Bjørn N. Freeman-Benson & Alan Borning (1992): *Integrating constraints with an object-oriented language*. In: *ECOOP*, pp. 268–286, doi:10.1007/BFb0053042.

[40] B.N. Freeman-Benson & A Borning (1992): *The design and implementation of* KALEIDOSCOPE*'90-A constraint imperative programming language*. In: *ICCL*, pp. 174–180, doi:10.1109/ICCL.1992.185480.

[41] Abdoulaye Gamatié (2010): *Designing embedded systems with the Signal programming language: synchronous, reactive specification*. Springer, doi:10.1007/978-1-4419-0941-1.

[42] Jeff Gennari, Shaun Hedrick, Fred Long, Justin Pincar & Robert Seacord (2007): *Ranged integers for the C programming language*. Technical Note CMU/SEI-2007-TN-027, Software Engineering Institute, Carnegie Mellon University. Available at `http://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=8265`.

[43] Yashdeep Godhal, Krishnendu Chatterjee & Thomas A Henzinger (2013): *Synthesis of AMBA AHB from formal specification: a case study*. *International Journal on Software Tools for Technology Transfer (STTT)* 15(5-6), pp. 585–601, doi:10.1007/s10009-011-0207-9.

[44] Nicolas Halbwachs (1993): *Synchronous programming of reactive systems*. *Engineering and Computer Science* 215, Springer, doi:10.1007/978-1-4757-2231-4. Available at `http://www-verimag.imag.fr/~halbwach/newbook.pdf`.

[45] Charles Antony Richard Hoare (1985, 2004): *Communicating sequential processes*. 178, Prentice-Hall. Available at `http://www.usingcsp.com/cspbook.pdf`.

[46] Gerard J. Holzmann: PROMELA *Language Reference (`http://spinroot.com/spin/Man/promela.html`)*. Available at `http://spinroot.com/spin/Man/promela.html`.

[47] Gerard J. Holzmann (2003): *The SPIN model checker: Primer and reference manual*. Addison-Wesley.

[48] Yong Jiang & Zongyan Qiu (2012): *S2N: model transformation from* SPIN *to* NUSMV. In: *SPIN*, pp. 255–260, doi:10.1007/978-3-642-31759-0_20.

[49] Barbara Jobstmann & Roderick Bloem (2006): *Optimizations for LTL synthesis*. In: *FMCAD*, pp. 117–124, doi:10.1109/FMCAD.2006.22.

[50] Barbara Jobstmann, Stefan Galler, Martin Weiglhofer & Roderick Bloem (2007): ANZU*: A tool for property synthesis*. In: *CAV*, pp. 258–262, doi:10.1007/978-3-540-73368-3_29.

[51] Barbara Jobstmann, Andreas Griesmayer & Roderick Bloem (2005): *Program repair as a game*. In: *CAV*, pp. 226–238, doi:10.1007/11513988_23.

[52] Muriel Jourdan, Fabienne Lagnier, R Maraninchi & Pascal Raymond (1994): *A multiparadigm language for reactive systems*. In: *ICCL*, pp. 211–218, doi:10.1109/ICCL.1994.288379.

[53] Robert M. Keller (1976): *Formal verification of parallel programs*. *CACM* 19(7), pp. 371–384, doi:10.1145/360248.360251.

[54] Yonit Kesten, Amir Pnueli & Li-on Raviv (1998): *Algorithmic verification of linear temporal logic specifications*. In: *ICALP*, 1443, pp. 1–16, doi:10.1007/BFb0055036.

[55] Uri Klein, Nir Piterman & Amir Pnueli (2012): *Effective synthesis of asynchronous systems from GR(1) specifications*. In: *VMCAI*, pp. 283–298, doi:10.1007/978-3-642-27940-9_19.

[56] M. Kloetzer & C. Belta (2008): *A fully automated framework for control of linear systems from temporal logic specifications*. *TAC* 53(1), pp. 287–297, doi:10.1109/TAC.2007.914952.

[57] H. Kress-Gazit, G.E. Fainekos & G.J. Pappas (2009): *Temporal-logic-based reactive mission and motion planning*. *IEEE Transactions on Robotics (TRO)* 25(6), pp. 1370–1381, doi:10.1109/TRO.2009.2030225.

[58] Daniel Kroening & Ofer Strichman (2008): *Decision procedures: An algorithmic point of view*. Springer.

[59] O. Kupferman & M.Y. Vardi (2005): *Safraless decision procedures*. In: *FOCS*, pp. 531–540, doi:10.1109/SFCS.2005.66.

[60] Orna Kupferman (2012): *Recent challenges and ideas in temporal synthesis*. In: *SOFSEM*, pp. 88–98, doi:10.1007/978-3-642-27660-6_8.

[61] Leslie Lamport (1994): *The Temporal Logic of Actions*. *ACM Trans. Program. Lang. Syst.* 16(3), pp. 872–923, doi:10.1145/177492.177726.

[62] Leslie Lamport (2002): *Specifying systems: The TLA+ language and tools or hardware and software engineers*. Addison-Wesley. Available at `http://research.microsoft.com/en-us/um/people/lamport/tla/book.html`.

[63] Leslie Lamport & Fred B. Schneider (1985): *Constraints: A uniform approach to aliasing and typing*. In: *POPL*, pp. 205–216, doi:10.1145/318593.318640.

[64] K Rustan M Leino (2010): *Dafny: An automatic program verifier for functional correctness*. In: *LPAR*, 6355, pp. 348–370, doi:10.1007/978-3-642-17511-4_20.

[65] A Solar Lezama (2008): *Program synthesis by sketching*. Ph.D. thesis, Citeseer. Available at `http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-177.html`.

[66] Orna Lichtenstein, Amir Pnueli & Lenore Zuck (1985): *The glory of the past*. In: *Logics of Programs*, 193, pp. 196–218, doi:10.1007/3-540-15648-8_16.

[67] Scott C. Livingston, Richard M. Murray & Joel W. Burdick (2012): *Backtracking temporal logic synthesis for uncertain environments*. In: *ICRA*, pp. 5163–5170, doi:10.1109/ICRA.2012.6225208.

[68] Gus Lopez, Bjorn Freeman-Benson & Alan Borning (1994): *Implementing constraint imperative programming languages: The* KALEIDOSCOPE*'93 virtual machine*. In: *OOPSLA*, pp. 259–271, doi:10.1145/191080.191118.

[69] Z Manna & Amir Pnueli (1990): *Tools and rules for the practicing verifier*. Technical Report CS-TR-90-1321, Stanford University, CA, USA. Available at `http://i.stanford.edu/pub/cstr/reports/cs/tr/90/1321/CS-TR-90-1321.pdf`.

[70] Zohar Manna & Amir Pnueli (1989): *The anchored version of the temporal framework*. In: *Linear time, branching time and partial order in Logics and models for concurrency*, *LNCS* 354, Springer, pp. 201–284, doi:10.1007/BFb0013024.

[71] Zohar Manna & Amir Pnueli (1990): *A hierarchy of temporal properties*. In: *PODC*, pp. 377–410, doi:10.1145/93385.93442.

[72] Shahar Maoz & Yaniv Sa'ar (2011): ASPECTLTL*: An aspect language for LTL specifications*. In: *Aspect-oriented Software Development (AOSD)*, ACM, pp. 19–30, doi:10.1145/1960275.1960280.

[73] John McCarthy (1959): *A basis for a mathematical theory of computation*. In P. Braffort & D. Hirschberg, editors: *Computer Programming and Formal Systems*, *Studies in Logic and the Foundations of Mathematics* 26, North-Holland, pp. 33–70, doi:10.1016/S0049-237X(09)70099-0.

[74] Kenneth Lauchlin McMillan (1992): *Symbolic model checking: An approach to the state explosion problem*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, USA, doi:10.1007/978-1-4615-3190-6. Available at `http://www.kenmcmil.com/pubs/thesis.pdf`. UMI Order No. GAX92-24209.

[75] George H Mealy (1955): *A method for synthesizing sequential circuits*. *Bell System Technical Journal* 34(5), pp. 1045–1079, doi:10.1002/j.1538-7305.1955.tb03788.x.

[76] Edward F Moore (1956): *Gedanken-experiments on sequential machines*. *Automata studies* 34, pp. 129–153.

[77] A. Morgenstern (2010): *Symbolic controller synthesis for LTL specifications*. Ph.D. thesis, Computer Science. Available at `http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:hbz:386-kluedo-25721`.

[78] A. Morgenstern & K. Schneider (2011): *A LTL fragment for GR(1)-synthesis*. *EPTCS* 50, pp. 33–45, doi:10.4204/EPTCS.50.3.

[79] David E Muller, Ahmed Saoudi & Paul E Schupp (1986): *Alternating automata, the weak monadic theory of the tree, and its complexity*. In: *ICALP*, pp. 275–283, doi:10.1007/3-540-16761-7_77.

[80] Elie Najm & Frank Olsen (1996): *Protocol verification with Reactive* PROELA*/RSPIN*. In: *SPIN*, pp. 109–128. Available at `http://spinroot.com/spin/Workshops/ws96/Ol.pdf`.

[81] Elie Najm & Frank Olsen (1996): *Reactive EFSMs — Reactive Promela/RSPIN*. In: *TACAS*, pp. 349–368, doi:10.1007/3-540-61042-1_54.

[82] Shipra Panda & Fabio Somenzi (1995): *Who are the variables in your neighbourhood*. In: *ICCAD*, pp. 74–77, doi:10.1109/ICCAD.1995.479994.

[83] Nir Piterman, Amir Pnueli & Yaniv Sa'ar (2006): *Synthesis of reactive(1) designs*. In: *VMCAI*, pp. 364–380, doi:10.1007/11609773_24.

[84] A. Pnueli & R. Rosner (1989): *On the synthesis of a reactive module*. In: *POPL*, pp. 179–190, doi:10.1145/75277.75293.

[85] Amir Pnueli (1977): *The temporal logic of programs*. In: *FOCS*, pp. 46 –57, doi:10.1109/SFCS.1977.32.

[86] Amir Pnueli & Uri Klein (2009): *Synthesis of programs from temporal property specifications*. In: *MEMOCODE*, pp. 1–7, doi:10.1109/MEMCOD.2009.5185372.

[87] Amir Pnueli & Roni Rosner (1989): *On the synthesis of an asynchronous reactive module*. In: *ICALP*, pp. 652–671, doi:10.1007/BFb0035790.

[88] Amir Pnueli, Yaniv Sa'ar & Lenore D Zuck (2010): JTLV*: A framework for developing verification algorithms*. In: *CAV*, pp. 171–174, doi:10.1007/978-3-642-14295-6_18.

[89] Roni Rosner (1992): *Modular synthesis of reactive systems*. Ph.D. thesis, Weizmann Institute of Science, Rehovot, Israel. Available at `http://www.researchgate.net/publication/238759536_Modular_synthesis_of_reactive_systems/file/50463527f8b648c3ba.pdf`.

[90] Richard Rudell (1993): *Dynamic variable ordering for ordered binary decision diagrams*. In: *ICCAD*, pp. 42–47, doi:10.1109/ICCAD.1993.580029.

[91] Matthias Schlaipfer, Georg Hofferek & Roderick Bloem (2012): *Generalized reactivity(1) synthesis without a monolithic strategy*. In: *HSVT*, pp. 20–34, doi:10.1007/978-3-642-34188-5_6.

[92] Klaus Schneider (2004): *Verification of reactive systems: formal methods and algorithms*. Springer, doi:10.1007/978-3-662-10778-2.

[93] Saqib Sohail, Fabio Somenzi & Kavita Ravi (2008): *A hybrid algorithm for LTL games*. In: *VMCAI*, pp. 309–323, doi:10.1007/978-3-540-78163-9_26.

[94] Fabio Somenzi (2012): CUDD*: CU Decision Diagram package - release 2.5.0*. University of Colorado at Boulder. Available at `http://vlsi.colorado.edu/~fabio/CUDD/cuddIntro.html`.

[95] Harald Søndergaard & Peter Sestoft (1992): *Non-determinism in functional languages*. *The Computer Journal* 35(5), pp. 514–523, doi:10.1093/comjnl/35.5.514.

[96] Wolfgang Thomas (2008): *Solution of Church's Problem: A tutorial*. *New Perspectives on Games and interaction* 5.

[97] Peter Van-Roy & Seif Haridi (2004): *Concepts, techniques, and models of computer programming*. MIT press.

[98] Moshe Y Vardi (1995): *Alternating automata and program verification*. In: *Computer Science Today*, Springer, pp. 471–485, doi:10.1007/BFb0015261.

[99] Moshe Y Vardi (1996): *An automata-theoretic approach to linear temporal logic*. In: *Logics for concurrency*, *LNCS* 1043, Springer, pp. 238–266, doi:10.1007/3-540-60915-6_6.

[100] Igor Walukiewicz (2004): *A Landscape with games in the background*. *LICS* 0, pp. 356–366, doi:10.1109/LICS.2004.1319630.

[101] Tichakorn Wongpiromsarn, Ufuk Topcu & Richard M Murray (2013): *Synthesis of control protocols for autonomous systems*. *Unmanned Systems* 1(01), pp. 21–39, doi:10.1142/S2301385013500027.